


ITFreeTraining




DDR Memory

For the free video please see
<http://itfreetraining.com/ap/3a40>

In this video from ITFreeTraining, I will look at DDR memory or Double Data Rate memory. DDR memory doubles the speed of data transfer from the previous memory type. DDR keeps being improved and is still the number one memory type sold on the market.

Double Data Rate



- Sends data on rise and fall of timing signal



0:16 Double Data Rate memory is different from previous memory modules in that it sends data on the rise and fall of the timing signal. The timing signal is generated by the memory controller and is sent to the memory module to control when commands and data are sent or received. As you can see, at the high and low point of the timing signal, data is sent. Before this, data was sent once for a full rise and fall of the timing signal. Essentially this method allows twice as much data to be sent and thus the name, Double Data Rate.

Double data rate has the obvious advantage that you are transferring more data per cycle. You could also increase the speed of the timing signal, for example by doubling the signal speed to get the same result. The problem with this approach is, as the timing signal gets faster, it is harder to control and it is subject to interference. For this reason, in the case of memory, you get a better result by increasing the data rate rather than the speed of the timing signal.

DDR1



1:19 The first standard of DDR otherwise was known as DDR1. DDR1 is not listed as a current exam objective, but I think it is good to look at it to give us an understanding of how far we have come.

DDR1 was released in 2000. There was other memory on the market that gave good performance, but DDR1 became the market leader due to its cheap price and good performance.

There are different versions of DDR1 - you will find they generally will operate in the 2.5 to 2.6 voltage range. Higher voltage means higher power consumption, but more importantly more heat. As components get smaller, the heat needs to be reduced, otherwise the components will fail.

The clock rate of DDR1 is between 100 and 200Mhz. This is the speed the memory controller communicates with the memory modules. Remember that DDR transfers twice as much data per clock cycle. For this reason, you will see different methods used to measure the speed of the memory other than the clock rate. More on that later in the video.

DDR1 has 184 pins. It requires a lot of pins as data is transferred in parallel. A lot of components in computers have converted to serial for communications. Memory Modules have not. Later in the video we will look at some of the reasons why this has not occurred.

DDR1 has a 2-bit prefetch. Prefetch is how much data, internally, the memory module can

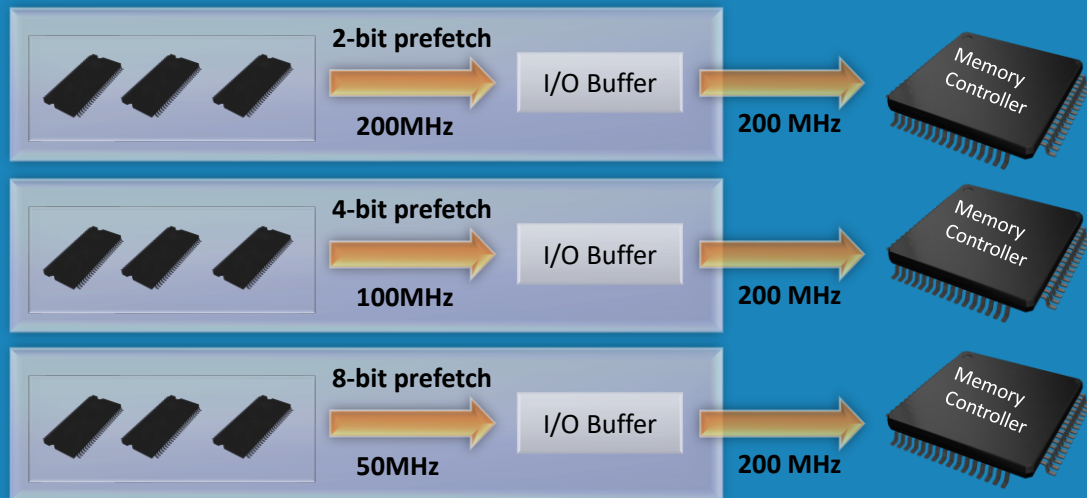
transfer to its buffers at once. The more data the memory module can put into its memory buffer, the faster it can transfer data to the memory controller.

The n in the measurement is the data width. Data width is essentially how much data can be transferred out of a single chip at once. The number is how much data the chip will get, at once, internally.

In the case of this chip, $2n$ means it can internally transfer twice n bits internally at once. To understand how this works, consider it like you are in a super market checkout. You realize you need some additional items so you go back into the store and get them. You have two hands so you can get two things. The checkout can only scan one item at a time, but the time taken to do this is small compared to how long it took to get the items.

So, in this example, the chip brings in two bits of data for every one bit of data it transfers out of the chip. Let's have a closer look at how having a bigger prefetch has on internal memory speed.

Prefetch



3:52 To understand how prefetch works, let's have a look at how it works in a memory module. In a memory module, consider you have DRAM chips. Usually eight on each side for a total of 16. When accessing data from the DRAM chips, it is stored in a buffer.

Once in this buffer, the data is transferred to the memory controller. If you consider our shopping market example, the buffer is like the conveyor belt leading to the cash register. If you can keep it full, you can maximum how many items go through the cash register. The higher the prefetch, the more fully you can keep the buffer. If the buffer runs out, no items are being transferred.

To understand better how this effects how much data we can transfer, consider the prefetch is 2 bits which is the case in DDR1. Let's also consider that the internal clock speed of the memory is 200MHz and also the bus speed to the memory controller is 200MHz.

Now let's also consider that we have a second memory module, however this memory module has a 4-bit prefetch. Once again, the bus speed to the memory controller is 200MHz.

Since the prefetch is twice the previous example, the internal bus speed only needs to be 100MHz or half the previous example.

If I now consider a third example, this time there is a prefetch of 8 bits. This is four times the first example and thus the internal bus only needs to be 50MHz.

These three examples will transfer the same amount of data. The only difference between each is the amount of prefetch each can do. As we will see shortly, newer designed DDR has a higher prefetch. This is one way the memory modules can transfer more data; however, they have lower internal clock speeds and bus speeds. You can start to understand why memory modules are still able to transfer in parallel while other devices in the computer are using serial. It's simply because, even though data transfer speeds have increased, the speed of the buses and internal speeds have only increased slowly. In the future, we may see a change to serial transmission, but at present parallel is working fine.

With all these different factors, it would be good to have a way of measuring speed. Back in DDR1 days there was a method commonly used. It is not used so much now, but good to know if you are working on an old computer.

PC (MB/Sec)



$PC \text{ (DDR Type)} = \text{Clock rate} * 2 \text{ transfer/sec} * 8$

$PC_1 = PC = 200 * 2 * 8$

$PC = 200 * 2 * 8 = 3200$

(PC1) One will always be dropped with DDR1. Making PC1 written as PC

6:24 To measure the speed of memory, the memory may have the speed in megabytes per second printed on the sticker of the memory module. If I consider my DDR1 example memory module, on the sticker is the measurement of how fast in megabytes per second the memory is.

To calculate this, the following formula is used. The clock rate is the bus speed to the memory module. Since double data rate transfers twice on each timing signal we need to multiply by 2. Lastly, 64 bits or 8 bytes are transferred at once so we need to multiply by 8.

To work out the speed for this memory, I need to substitute in the clock rate in the formula. Note that with DDR1, the one is dropped so it is just PC not PC1. This is one way you can identify memory modules by looking at what follows after PC on the sticker.

In this example, this gives us a result of 3200. Essentially this means the memory module will transfer at a maximum speed of 3200 megabytes per second. In the real world, this is a maximum speed due to a lot of other factors, so you won't get this speed consistently.

The PC measurement may not always be on the memory sticker, particular with the new DDR memory modules. If this is the case, there is another measurement of the speed on the memory module.

DDR Speed



DDR (DDR Type) = Clock rate * 2

~~DDR1~~ = DDR = 200 * 2

DDR = 200 * 2 = 400

(DDR1) One will most of the time be dropped with DDR1. Making DDR1 written as DDR


7:48 Looking once again at our example memory module this particular memory module has both of the measurements used to measure what performance we can expect to get. Some memory modules may have only one. You can see after DDR there are some numbers.

This number is calculated by multiplying the clock rate by two. Like before, a number is added to indicate which version of DDR RAM is being used. Most of the time it will be dropped for DDR1. In other DDR memory, the number will be added so people know which version of DDR they are looking at.

I will put the clock rate in my formula and will get the result of 400. Essentially, since double data rate can transfer twice per cycle, this is why we multiply the clock rate by two. Keep in mind this is a measure of the clock rate not the amount of data transferred. To get the amount of data transferred multiply this number by eight and we will get 3200 which is the figure after PC which I looked at in the previous example. When you compare one memory module to another, make sure you compare the same measurements or convert one to the other by multiplying or dividing by eight.

Now that we understand DDR1, I will have a look at DDR2.

DDR2



9:03 DDR2 was released in 2003. There were other memory modules on the market, but due to its low cost compared with the others, it has become the dominant memory on the market. DDR2 reduced the voltage down to 1.8 volts. Lower voltage means less heat, which allows for more transistors on a chip. This becomes more important when trying to increase the amount of data that can be stored on a memory module.

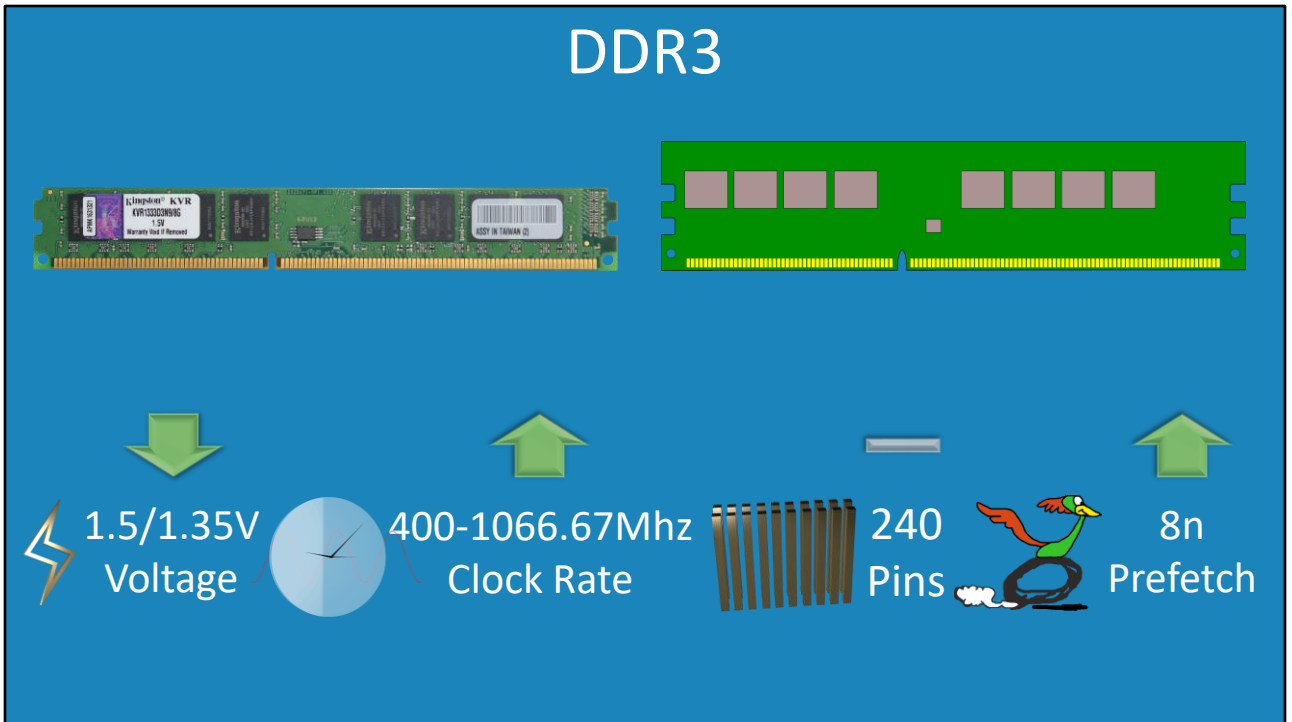
The next change was that the clock rate was increased. The clock rate has essentially doubled from DDR1. The faster clock rate allows more data to be transferred over the bus to the memory controller.

The next change was that the pin count increased to 240 from 184. The memory module still transfers 64bits at once, and the additional pins allow for better control of the memory module. With the extra pins, this means DDR2 is not compatible with DDR1.

The next change is the prefetch was increased from two to four. This change essentially means that the internal clock rate of the memory module only needs to be a quarter of the clock rate. One way you can improve performance is to increase the amount of data sent internally, which allows you to lower the internal clock rate. Having a lower internal clock rate means you have less timing problems. Not to be confused with the clock rate to the memory controller, the internal clock rate is the clock rate used to access the chips on the memory module, not the memory module itself.

As computers get faster, it gets harder and harder to control them. Little tricks like these allow the memory to keep up with the rest of the computer, even though in reality its speed is not the same as the rest of the computer.

DDR3



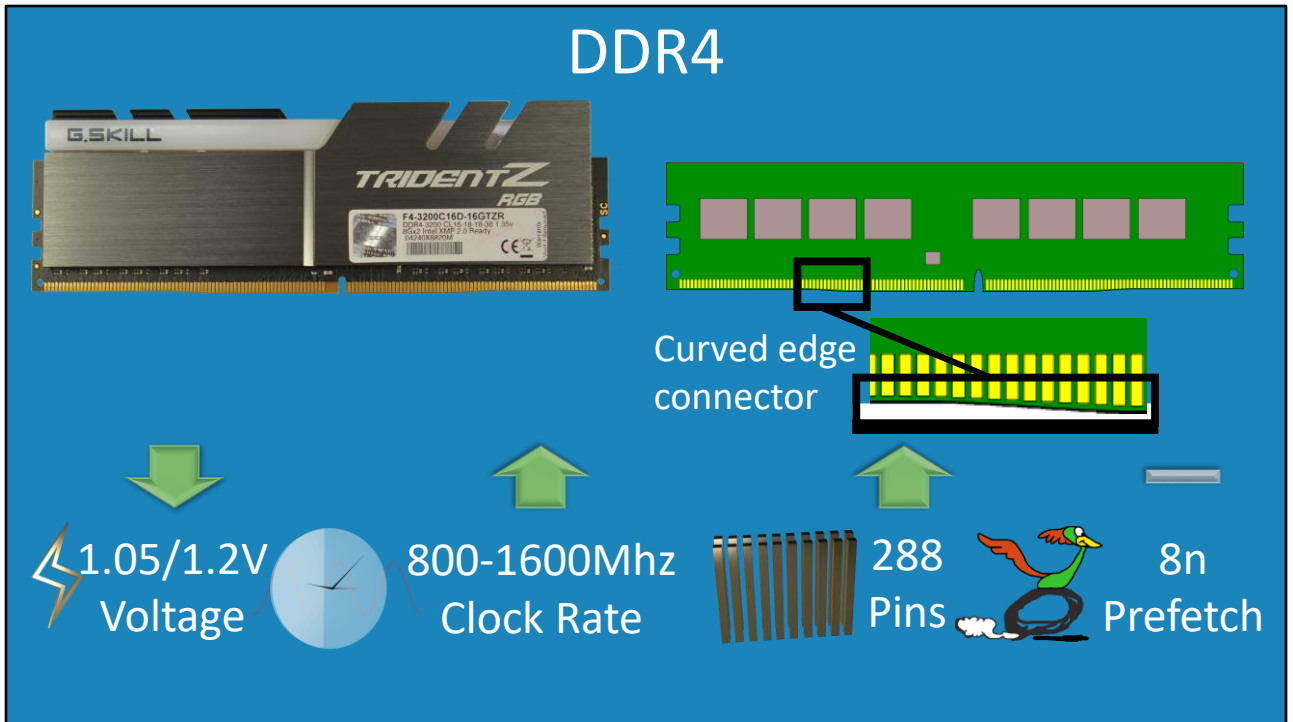
10:49 The next memory I will look at is DDR3. You will notice in this example the memory module is very low in height compared with the others ones. The memory has to be the same length, but the manufacturer has some freedom in the height they make the memory module. As long as it has the correct number of pins and notch in the right location, it will work. Each different standard of DDR has a small notch around the center of the memory module and each DDR type has the notch in a slightly different place, preventing the wrong memory being put in the wrong slots. At the end of the video I will compare all the different standards of DDR and then this will make more sense.

Like the memory before it, the voltage has once again been reduced. This helps reduce the heat generated by the memory module. As before, the clock rate has also increased. As computers get faster, the memory needs to get faster, as well. Remember, this clock rate refers to the bus between the memory module and the memory controller.

With DDR3, the number of pins remains the same as DDR2. Even though the pins are the same, remember the notch in the memory prevents the memory being put in the wrong way.

Lastly, the prefetch has been increased from four to eight. With the increase in prefetch, this means the internal clock rate of the memory module does not have to increase. Since the prefetch has increased with each standard of DDR so far, the internal clock rate of the memory module has remained about the same.

DDR4







12:22 The last memory that I will look at is DDR4. DDR4 was released in 2014. You will notice in this example that the memory module has a heat sync attached to the memory. You will find most DDR4 will have a heat sync, some DDR3 will also. It is up to the manufacturer to decide if they want to add one or not. You will find that this is more common practice as memory modules pack more transistors onto the chip as they increase the capacity. More transistors mean more heat.


With DDR4, the voltage has decreased to between 1.05 to 1.2 volts. Decreasing the voltage becomes important as the clock rate has increased from 800 to 1600Mhz. The increased clock rate also increases the amount of heat the memory module creates.

The next change is the number of pins has increased to 288 pins. With the increase in the number of pins, the edge of the connector is slightly curved. This is done so that the memory module can be put in the slot with less force.

Like DDR3, DDR4 uses 8-bit prefetch. Even though this has not increased, since the clock rate has increased, the overall speed of the memory module has increased. At the time this video was created, DDR5 was not yet on the market. DDR5 once again promises to increase the speed and performance of memory.

Summary

	Voltage 	Clock Rate 	Pins 	Prefetch 
DDR1	2.5/2.6	100-200	184	2n
DDR2	1.8	200-533.33	240	4n
DDR3	1.35/1.50	400-1066.67	240	8n
DDR4	1.05/1.2	800-1600	288	8n



13:49 To end this video, I will perform a quick summary of all the different DDR types. You can see that with each newer edition of the DDR standard, the voltage has decreased for each specification. Notice that the clock rate has increased for each different specification of DDR. This helps make the memory faster; however, a big difference that makes the memory faster is also the amount of prefetch each specification has. A larger prefetch allows more data in the memory module to be accessed at once. This also allows the internal clock of the memory module to be slower, which therefore increases the stability of the memory.

Although each DDR is the same size, some have a different number of pins. However, you will notice the notch on each different specification is in a different location. This means you won't be able to put the wrong memory in the wrong slot and also it can only go in one way. This is why you should never force DDR memory in when you are installing it. If you find it will not go in, you most likely have the memory module turned around the wrong direction.

That concludes this video on DDR memory. I hope you have found this video useful. Until the next video from us, I would like to thank you for watching.

References

<https://www.hardwaresecrets.com/everything-you-need-to-know-about-ddr-ddr2-and-ddr3-memories/5/> "Everything You Need To Know About DDR, DDR2 and DDR3 Memories"

Credits

Trainer: Austin Mason <http://ITFreeTraining.com>

Voice Talent: HP Lewis <http://hplewis.com>

Quality Assurance: Brett Batson <http://www.pbb-proofreading.uk>